PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04Q

A2

(11) International Publication Number: WO 98/26611

(43) International Publication Date: 18 June 1998 (18.06.98)

(21) International Application Number: PCT/US97/22456

(22) International Filing Date: 9 December 1997 (09.12.97)

(30) Priority Data:
08/766,423 12 December 1996 (12.12.96) US

(71) Applicant: CASCADE COMMUNICATIONS CORP. [US/US]; 5 Carlisle Road, Westford, MA 01886 (US).

(72) Inventor: GANMUKHI, Mahesh, N.; 1286 Curve Street, Carlisle, MA 01741 (US).

(74) Agents: LEBOVICI, Victor, B. et al.; Weingarten, Schurgin, Gagnebin & Hayes LLP, Ten Post Office Square, Boston, MA 02109 (US).

(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published
Without international search report and to be republished upon receipt of that report.

(54) Title: SWITCH FABRIC SWITCHOVER IN AN ATM NETWORK SWITCH

(57) Abstract

A method and apparatus for performing switchover in a network switch having first and second switch fabrics from one of the switch fabrics to the other of the switch fabrics so as to avoid cell loss during the switchover process. One of the switch fabrics is active and the other switch fabric is in a standby mode. In one embodiment switchover occurs immediately following detection of conditions indicating that the active switch fabric has fully drained. In another embodiment switchover occurs following a timeout which assures full drainage of the active switch fabric under worst case conditions.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

TITLE OF THE INVENTION

SWITCH FABRIC SWITCHOVER IN AN ATM NETWORK SWITCH

FIELD OF THE INVENTION

The present invention relates to telecommunications and more particularly to a method and apparatus for performing switchover from one switch fabric to another switch fabric in a network switch.

BACKGROUND OF THE INVENTION

Network switches, and particularly Asynchronous Transfer Mode (ATM) network switches are frequently appearing in telecommunications networks managed by data network service providers to efficiently handle the voluminous voice and data traffic within such networks. With the increasing utilization of such network switches, concerns have arisen regarding the ability to maintain the network switches in operation following a component failure or while maintenance and/or repair are performed so as to avoid service interruption and undesirable loss of cells transmitted by subscribers.

One approach for performing switch fabric switchover in an ATM network switch is described in U.S. Patent No. 5,488,606. In the system therein disclosed redundant switch fabrics are provided. A cell marking technique is employed to mark received cells as being destined either for the master switch fabric or a backup switch fabric. One switch fabric is designated as a master and a second switch fabric is designated as a slave. The designations may be selectably swapped. Cells tagged with a master indication are diverted to the switch fabric designated as the master unit and cells tagged with a slave indication are diverted to the switch fabric designated as the slave. A cell reception unit is provided in each of the output modules and receives the cells from the switch fabric designated as the master at any given time. Upon an indication of a switchover condition, the switch fabric designated as the master unit is redesignated as the slave and the switch fabric designated as the slave is redesignated as the master. The cell reception units

continue to process cells which are tagged as "master" cells. Cells designated as master cells following the switchover condition are processed by the newly designated master switch fabric.

SUMMARY OF THE INVENTION

In accordance with the present invention, a network switch is provided which permits rapid switchover from a first switch fabric to a second switch fabric upon detection of a switchover condition. The network switch includes a plurality of input modules for receiving ATM cells over respective ATM links, a plurality of output modules for transmitting ATM cells over ATM links and redundant selectable switch fabrics. One of the switch fabrics is operative at any given time and the other switch fabric remains in a standby mode. The active switch fabric selectively forwards cells received at one of the input modules to one or more of the output modules for transmission onto an associated ATM link. Input buffers provided in each of the network switch input modules are appropriately sized to accommodate an incoming cell stream when backpressure is applied to the input modules as a consequence of congestion in the switch fabric or the initiation of a switchover condition. Output buffers are provided in each of the network switch output modules to buffer cells received from the switch fabric prior to transmission over the respective ATM link.

In the event of the detection of a switchover condition the transmission of cells from the input modules to the output modules is stopped in a controlled fashion. More specifically, following receipt of a signal indicating that a switchover condition has occurred, the transmission of ATM cells which are in progress from the input module to the switch fabric is completed. Cells which have not commenced transmission from the input module to the switch fabric are prevented from initiating transit.

Detection circuitry is provided in each of the output modules to detect when all cells destined for the respective port have been drained from the switch fabric. When all cells have been drained from the switch fabric such condition is detected by a switch processor. In response to the detection of a signal indicating that all cells have been drained from the switch fabric, the active switch fabric is converted to the

inactive or standby switch fabric and the standby switch fabric is activated so as to become the active switch fabric. Thereafter transmission of cells through the newly activated switch fabric is permitted to continue.

BRIEF DESCRIPTION OF THE DRAWING

The invention will be more fully understood by reference to the Detailed Description of the Invention in conjunction with the Drawing of which:

Fig. 1 is a block diagram of a network switch incorporating redundant switch fabrics in accordance with the present invention;

Fig. 2. is a timing diagram illustrating signals employed to achieve switchover in accordance with the presently disclosed technique; and

Fig. 3 is a flow diagram illustrating the method for accomplishing switchover in accordance with the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Referring to Fig. 1, a network switch 10 includes a plurality of input modules 12 a plurality of output modules 14 and switch fabrics 16a and 16b for selectively forwarding data units received at inputs of respective input modules 12 to one or more of the output modules 14. Although associated input modules 12 and the output modules 14 are typically fabricated on a single printed circuit card, they are depicted separately in Fig. 1 for clarity. The input modules 12 receive streams of data units such as Asynchronous Transfer Mode (ATM) cells, frames or packets over communication links 20. Each of the input modules 12 includes buffers 22 for storage of received data units. Similarly, the output modules 14 each includes buffers 24 which store data units prior to transmittal over respective output links 26.

Redundant switch fabrics 16a and 16b are provided so as to permit switchover from one switch fabric to the other of the switch fabrics in the event that some abnormality in the operation of a switch fabric is detected or in the event that it is desired to remove a switch fabric for purposes of testing or maintenance. The presently disclosed apparatus and method provides for switch fabric switchover to be accomplished so as to avoid the loss of data units, such as ATM cells, frames or

packets, during the switchover process.

Switch fabric switchover is controlled by a switch processor 28. Upon detection of a condition warranting switchover from one switch fabric, such as switch fabric 16a, to another switch fabric, such as switch fabric 16b, the switch processor 28 broadcasts a signal (SYNC_TIME_L) to all input modules 12. This signal provides an indication to all input and output modules that switchover is proceeding and that all input modules should cease transmission of data units through the switch fabric following the completion of the transmission of any data units then in transit.

More specifically, referring to Fig. 1 in conjunction with Fig. 2 the SYNC_TIME_L signal is provided to both the input modules 12 and the output modules 14 at time T_0 as illustrated in Fig. 2. Upon receipt of the SYNC_TIME_L signal, the input modules 12 complete forwarding to the switch fabrics, any data units that are then in transit. However, the input modules cease forwarding any additional data units through the operative switch fabric. Additional data units in the input modules which are not in transit to the switch fabrics and data units received by an input module after receipt of the SYNC_TIME_L signal are instead accumulated in data buffers 22 associated with the respective input modules 12 pending completion of the switchover process. As will be apparent, the size of the data buffers 22 are dependent upon the rapidity of the switchover process, i.e., the more rapid the switchover process, the smaller the size of required input modules data buffers 22.

In practice, a certain amount of time will be required following the cessation of transmission of data units to the switch fabric to assure that the switch fabric has fully drained of data units since the switch fabric itself may contain data storage buffers for storage of data units in transit through the switch fabric. In accordance with the presently disclosed technique, it is desirable not to proceed with switchover until all data units have drained from the switch fabric to avoid loss of data units since corruption and/or loss of data units will in many cases necessitate retransmission of the effected data units and increase network traffic. It should be appreciated that loss of some data units may be inevitable since corruption of data units may result from switch fabric failures which prompted the switchover condition

initially.

In one embodiment, respective output modules 14 detect a condition which indicates that the operative switch fabric has drained of valid data units received at the input modules over respective communication links. Each output module 14 provides an indication that it is no longer receiving valid data units. When the switch processor 28 detects that all output modules are no longer receiving valid data units, the switch processor 28 provides a signal to the switch fabric causing switchover from one of the switch fabrics 16a and 16b to the other of the switch fabrics to occur. Following the deactivation of the active switch fabric and the activation of the standby switch fabric, the input modules 12 continue to forward data units to the newly activated and operative switch fabric for delivery of such data units to selected output modules 14. By sensing that the operative switch fabric has fully drained and permitting switchover to proceed as soon as this condition has been detected, loss of data units due to the switchover process is avoided and the size of data buffers 22 which are needed to accumulate data units during the period that data unit forwarding has been halted is reduced. Since the size of data buffers 22 can significantly contribute to the cost of a network switch, it is desirable to minimize the size and consequently, the cost of such buffers.

More specifically, upon detection of a switchover condition, such as a defective switch fabric or a user input condition indicating a desire to change from one switch fabric to another switch fabric, the switch processor forwards the SYNC_TIME_L signal (see Fig. 2) to input modules over signal line 30 and to output modules over signal line 32. Depending upon the particular implementation, it will be understood that signal lines 30 and 32 may be combined as a single signal line. The input modules 12, upon detection of the SYNC_TIME_L signal, complete the transmission to the operative switch fabric of any data units, such as ATM cells, already in transit. The input modules 12, however, do not forward any additional data units. Instead, the input modules store data units not in transit to the switch fabric in data buffers 22 associated with the respective input modules and additionally, store data units received after receipt of the SYNC_TIME_L signal in the respective data buffers 22.

The operative switch fabric, 16a or 16b, selectively forwards the data units received from the various input modules to one or more of the output modules as is known in the art.

The output modules 14 continue to monitor the received data streams from the switch fabric for, among other things, a predetermined character. More particularly, in a preferred embodiment comprising an ATM network switch, the output modules monitor the input data streams for an idle cell. Upon receipt of a predetermined number of such idle cells in succession, the output modules provide an output signal identified as the NREC signal as illustrated in Fig. 2 over signal line 34. The NREC signal may comprise a wired AND signal. For example, the driver for the NREC signal within the output modules may comprise an open collector npn transistor with its emitter coupled to ground and the collector coupled to the signal line 34. A common pull-up resistor is coupled to the signal line 34 with the open collectors comprising the outputs coupled to the single pull-up resistor. Thus, the NREC signal will go high only when none of the output modules npn drive transistors are on. Accordingly, when all of the drive transistors are biased to an off condition, upon detection by all of the output modules of the receipt of at least a predetermined number of idle cells in succession following the assertion of the SYNC_TIME_L signal, the NREC signal on signal line 34 is asserted as indicated at time T_1 in Fig. 2. Otherwise stated, the NREC signal is asserted when the last of the output modules has recognized that it has received a predetermined number of successive idle cells.

More specifically, following the assertion of the SYNC_TIME_L signal, a counter within each of the output modules is incremented each time an idle cell is received. When "n" back to back idle cells are received, it is assumed that the switch fabric has no further cells to forward. During the idle cell counting period, if a non-idle cell is received, the counter is reset and the count begins again. In a preferred embodiment of the invention, "n" is set to 16 which corresponds to a time larger than the worst case propagation delay for a cell through the switch fabric.

The NREC signal is coupled to the switch processor 28. While the present implementation has been illustrated as a wired AND signal, it should be appreciated

that each of the output modules could forward a separate signal to the switch processor 28 with the discrimination occurring at the switch processor. Additionally, any other suitable technique may be employed for communicating to the switch processor 28 that each of the output modules has received at least a predetermined number of idle cells in succession.

The detection by the switch processor 28 of the fact that the NREC signal has been asserted provides an indication that the then operative switch fabric 16a or 16b has been fully drained of data units and accordingly, switch fabric switchover may proceed without loss of data units due to the switchover process. In response to detection of the assertion of the NREC signal at time T_1, the switch processor causes the SFA_L signal 29 to change state. The SFA_L signal is coupled to the switch fabrics 16a and 16b and constitutes the selection signal which causes one switch fabric to assume an active state and the other switch fabric to assume an inactive state. More specifically, when the SFA_L signal is low, the outputs of switch fabric A are coupled to the inputs of the respective output modules and the outputs of switch fabric B are decoupled from the inputs of the respective output modules. Further, when the SFA_L signal is high, the outputs of switch fabric A are decoupled from the inputs of the respective output modules and the outputs of switch fabric B are coupled to the inputs of the respective output modules. Such may be accomplished via a tri-stated bus, a multiplexing technique in which separate data lines are multiplexed to provide a selectable single output which is applied to the output modules or via any other suitable technique known in the art.

After the switch processor 28 causes the change of state of the SFA_L signal at time T_2 as illustrated in Fig. 2, the switch processor 28 causes the SYNC_TIME_L signal to go high at time T_3 thereby permitting the input modules to recommence the forwarding of data units, such as ATM cells, through the operative switch fabric to the output modules. Further, in response to the deassertion of the SYNC_TIME_L signal at time T_3 the output modules each cease to activate the NREC signal 34 (since the assertion of SYNC_TIME_L is a prerequisite to the assertion of NREC) and the NREC signal goes low thereby completing the switchover operation.

It should further be noted that it is necessary to account for a broken or

"babbling" switch fabric which, due to the nature of the switch fabric failure, may not be able to transmit proper idle cells. In such event, a timeout occurs following the passage of a predetermined time period following the assertion of SYNC_TIME_L. This period corresponds to the worst case period for complete drainage of the switch fabric.

The above described method is further illustrated in Fig. 3. Referring to Fig. 3, in decision step 40 a test is performed to determine if a switchover condition has been detected. If a switchover condition has been detected, as illustrated in step 42, the transmission of data units already in transit to the active switch fabric is completed and subsequent data units are stored in the data buffers of respective input modules. As indicated in step 44, following detection of the switchover condition, idle cell counters are preset in each of the output modules. As illustrated in decision step 46, each received cell is tested at the respective output modules to ascertain whether the received cell is an idle cell. If the received cell is an idle cell, the limit counter is decremented for the respective output module as illustrated in step 50. If the received cell is not an idle cell, the limit counter in the respective output module is preset and control is again passed to decision step 46 to ascertain whether the next received cell is an idle cell. Following step 50, as illustrated in step 52, the limit counter is tested to determine whether at least a predetermined number of successive idle cells have been received in each of the output modules. If a predetermined number of successive idle cells have been received in all output modules, as illustrated in step 54, control passes to step 56. As depicted in step 56 the active switch fabric is deactivated and the standby switch is activated. Following the transfer of control from the former active switch fabric to the former standby switch fabric, the switchover condition signal (SYNC_TIME_L) is reset as depicted in step 58. Once the switchover condition signal is reset, data forwarding through the newly activated switch fabric commences as illustrated in step 60. Thereafter, as illustrated in step 40, the system continues to monitor the switchover condition signal for a switchover condition signal. In the foregoing manner, control is passed smoothly from one of the first and second switch fabrics to the other switch fabric without cell loss due to the switchover process.

In another embodiment of the invention, the NREC signal is not employed and the complexity associated with the detection of a predetermined number of successive idle cells is avoided at the expense of enlarged data buffers 22. In this embodiment, after the assertion of the SYNC_TIME_L signal, the switch processor 28 waits for the worst case time period to assure full drainage of all data units from the switch fabric before initiating switchover.

In the event that the worst case wait technique is employed to assure full drainage of the switch fabric, the SA_L signal is driven to its opposite state to accomplish switchover following the passage of the predetermined time interval to accommodate the worst case switch fabric drainage period. As a consequence of the need to store incoming data units following assertion of the SYNC_TIME_L signal and prior to switch fabric switchover, when the worst case wait technique is employed greater input buffering is needed to avoid loss of incoming data units due to data overruns. Generally, the traffic through the switch fabric is not in a "many to one" traffic format in which a number of data units from various input modules are destined for a single output module and consequently, the worst case wait technique for assuring data unit drainage is disadvantageous for at least two reasons. First, the worst case wait technique requires larger data buffers 22 to be employed to store incoming data units during the prolonged period allocated for switch fabric drainage. Additionally, an additional time interval is interposed between the assertion of the SYNC_TIME_L signal and the actual switchover from the first switch fabric to the second switch fabric.

It will be appreciated that variations and modifications of the above described method and apparatus for performing switchover in a network switch will be apparent to those of ordinary skill in the art without departing from the inventive concepts described herein. Accordingly, the preferred embodiments described herein are to be viewed merely as illustrative, and not limiting, and the invention is to be limited solely by the scope and spirit of the appended claims.

CLAIMS

What is claimed is:

1. A method for deactivating a first switch fabric in a network switch having first and second switch fabrics and for activating a second switch fabric in response to an indication of a switchover condition, wherein said network switch has a plurality of input modules for receiving data units over an corresponding input communications link, a plurality of output modules for forwarding data units received from said input communications links to an associated output communications links and wherein one of said first and second switch fabrics is operative to forward data units received at said input modules to selected ones of said output modules, said method comprising the steps of:

in a first detecting step, detecting an indication of a switchover condition and in response to said detection of said switchover condition discontinuing the forwarding of data units received at said plurality of input modules to said first switch fabric;

in a second detecting step, following detection of said indication of a switchover condition, at each of said output modules, detecting the absence of received data units being forwarded from said first switch fabric for a predetermined period and providing an indication of such absence of received data units to a network switch processor;

in a third detecting step, detecting at said network switch processor the absence of received data units at all of said output modules for at least a predetermined period and, in response thereto deactivating said first switch fabric and activating said second switch fabric; and

following said third detecting step, forwarding data units stored in respective data buffers of said input modules to said second switch fabric for transmittal to selected ones of said plurality of output modules.

2. The method of claim 1 wherein said second detecting step of detecting the absence of a received data units comprises the step of detecting receipt of at least a predetermined number of successive idle cells at each of said output modules.

3. The method of claim 1 wherein said data units comprise fixed length data units.

4. The method of claim 1 wherein said fixed length data units comprise Asynchronous Transfer Mode (ATM) cells.

5. The method of claim 4 wherein in said second detecting step, said step of detecting the absence of a received data units comprises the step of detecting at least a predetermined number of idle cells at each of said output modules.

6. The method of claim 5 wherein said step of detecting at said network switch processor the absence of received data units at all of said output modules for at least a predetermined period comprises detecting at said network switch at least one signal which indicates that at least a predetermined number of idle cells have been received at each of said output modules following the detection of said indication of said switchover condition.

7. The method of claim 1 wherein said step of detecting a switchover condition comprises the step of detecting an abnormality in the operation said first switch fabric.

8. The method of claim 1 wherein said step of detecting a switchover condition comprises one of detecting an abnormality in the operation of said first switch fabric and detecting an indication from a user that switch fabric switchover should occur.

9. A method for deactivating a first switch fabric in a network switch having first and second switch fabrics and for activating a second switch fabric in response to an indication of a switchover condition, wherein said network switch has a plurality of input modules for receiving data units over an corresponding input communications link, a plurality of output modules for forwarding data units received from said input communications links to one or more of said output communications

links and wherein one of said first and second switch fabrics is operative to forward data units, said method comprising the steps of:

detecting an indication of a switchover condition and, in response to said detection;

discontinuing the forwarding of data units received at said plurality of input modules to said first switch fabric, and

storing in data buffers associated with each of said input modules, data units received at respective input modules following detection of said switchover condition;

detecting a predetermined event following the detection of said switchover condition and in response thereto;

deactivating said first switch fabric;

activating said second switch fabric; and

initiating the forwarding of data units through said second switch fabric.

10. The method of claim 9 wherein said step of detecting a predetermined event comprises the step of detecting the passage of a predetermined time interval.

11. The method of claim 10 wherein said step of detecting the passage of a predetermined time interval comprises the step of detecting the passage of a time interval corresponding to the time required for the forwarding of a predetermined number of data units from at least one of said input modules to at least one of said output modules.

12. The method of claim 11 wherein said data units comprise ATM cells.

13. A network switch comprising:

a plurality of input modules, each input module being operative to receive data units over a respective communications link;

a plurality of output modules;

a plurality of data buffers associated with respective input modules for storing data units received by the respective input modules;

a switch processor being operative to provide a control signal having at least a first and second state;

at least first and second switch fabrics,

said first and second switch fabrics being each being operative to forward data units received at said plurality of input modules to selected ones of said plurality of said output modules when said control signal is in said first and second state respectively and to discontinue the forwarding of data units received at said plurality of input modules to selected ones of said plurality of output modules when said control signal is in said second and first state respectively, wherein one of said first and second switch fabrics is operative when said control signal is in one of said first and second states and the other of said first and second switch fabrics is operative when said control signal is in the other of said at least first and second states;

said input modules being operative in response to a signal indicative of a switchover condition to cease transmission on data units from said input modules to said output modules through the operative one of said switch fabrics and to accumulate in the respective data buffers, data units received after recognition of said signal indicative of said switchover condition;

said output modules each being operative in response to said signal indicative of a switchover condition to activate a first character recognition signal upon receipt from the operative switch fabric of predetermined number of successive data units having a first predetermined value;

said switch processor being responsive to said first character recognition signals and an indication that all of said output modules have received at least said predetermined number of data units having said first predetermined value to cause said control signal to change from one of said first and second states to the other of said first and second states; and

said switch processor being further operative to cause said input modules to reinitiate forwarding of data units accumulated in respective data buffers of said input

modules following the change in state of said control signal so as to complete transfer of switch fabric operation from one of said at least first and second switch fabrics to the other of said at least first and second switch fabrics.

14. The network switch of claim 13 wherein said data units each have a fixed length.

15. The network switch of claim 14 wherein said data units comprise asynchronous transfer mode (ATM) cells.

16. The network switch of claim 15 wherein said input modules are operative to complete the forwarding through the operative switch fabric of any ATM cells in transit upon receipt of said signal indicative of a switchover condition and to cease forwarding of cells received at respective input modules but not yet forwarded to said operative switch fabric upon receipt of said signal indicative of said switchover condition.

17. The network switch of claim 15 wherein said data units data units having said first predetermined value comprise idle cells.

18. A network switch comprising:

a plurality of input modules, each input module being operative to receive data units over a respective communications link;

a plurality of output modules;

a plurality of data buffers associated with respective input modules for storing data units received by the respective input modules;

a switch processor being operative to provide a control signal having at least a first and second state;

at least first and second switch fabrics;

said first and second switch fabrics each being capable of forwarding data

units received at said plurality of input modules to selected ones of said plurality of said output modules when said control signal is in said first and second state respectively and to discontinue the forwarding of data units received at said plurality of input modules to selected ones of said plurality of output modules when said control signal is in said second and first state respectively, wherein one of said first and second switch fabrics is operative when said control signal is in one of said first and second states and the other of said first and second switch fabrics is operative when said control signal is in the other of said at least first and second states;

said input modules being operative in response to a signal indicative of a switchover condition to cease transmission on data units from said input modules to said output modules through the operative one of said switch fabrics and to accumulate in the respective data buffers, data units received at the respective input I.O modules after recognition of said signal indicative of said switchover condition;

said output modules each being operative in response to said signal indicative of a switchover condition to activate a first character recognition signal upon receipt from the operative switch fabric of predetermined number of successive data units having a first predetermined value;

said switch processor being operative in response to the passage of a predetermined time period following receipt of said signal indicative of said switch second states to the other of said first and second states; and

said switch processor being further operative to cause said input modules to reinitiate forwarding of data units accumulated in respective data buffers of said input modules following the change in state of said control signal.

19. The network switch of claim 18 wherein said data units each have a fixed length.

20. The network switch of claim 14 wherein said data units comprise asynchronous transfer mode (ATM) cells.

21. The network switch of claim 20 wherein said input modules are operative to

complete the forwarding through the respective switch fabric of any ATM cells in transit upon receipt of said signal indicative of a switchover condition and to cease forwarding of cells received at respective input modules but not yet forwarded to said operative switch fabric upon receipt of said signal indicative of said switchover condition.

22. The network switch of claim 18 wherein said predetermined time period corresponds to a time period specified to assure that all data units which have been forwarded from respective input modules have drained from said operative one of said at least first and second switch fabrics.

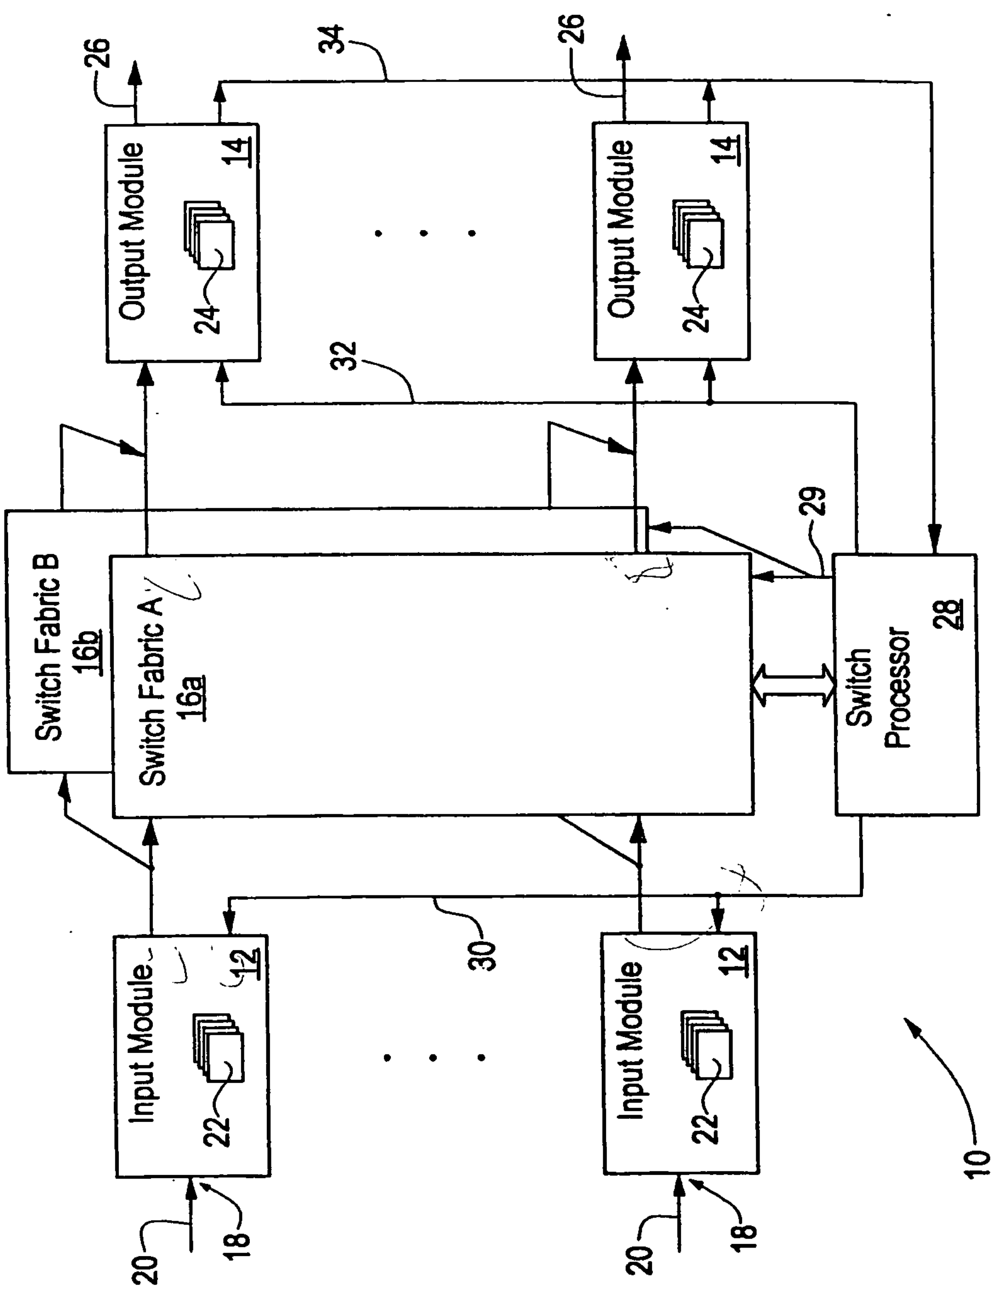
Output Module
14
24
26
Output Module
14
24
26
34
32
Switch Fabric B
16b
Switch Fabric A
16a
Switch
Processor
28
29
30
Input Module
12
22
Input Module
12
22
20
18
20
18
10

FIG. 1

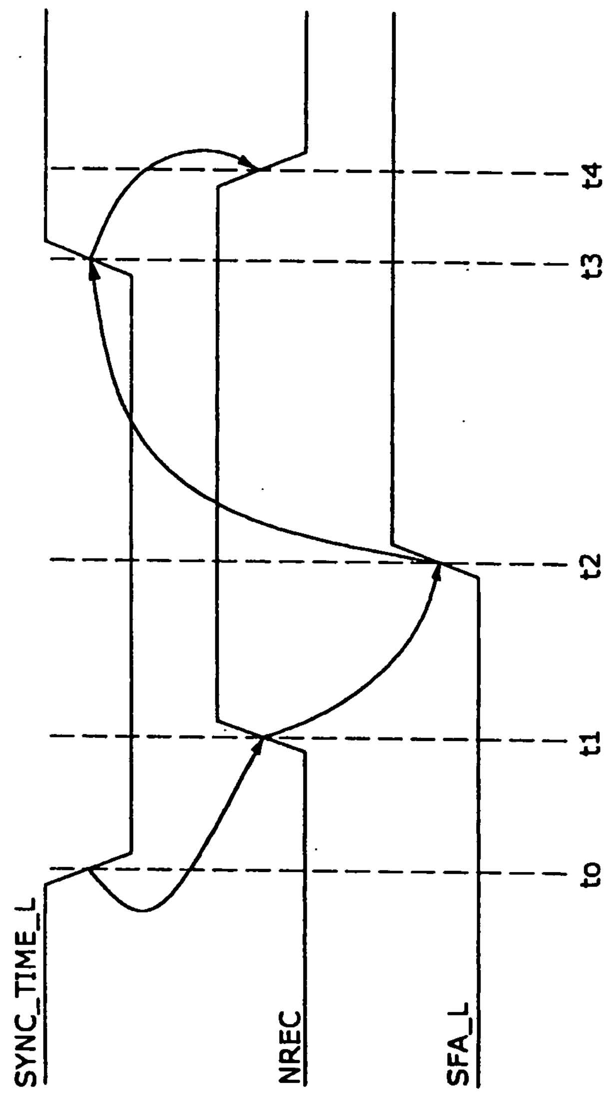

FIG. 2

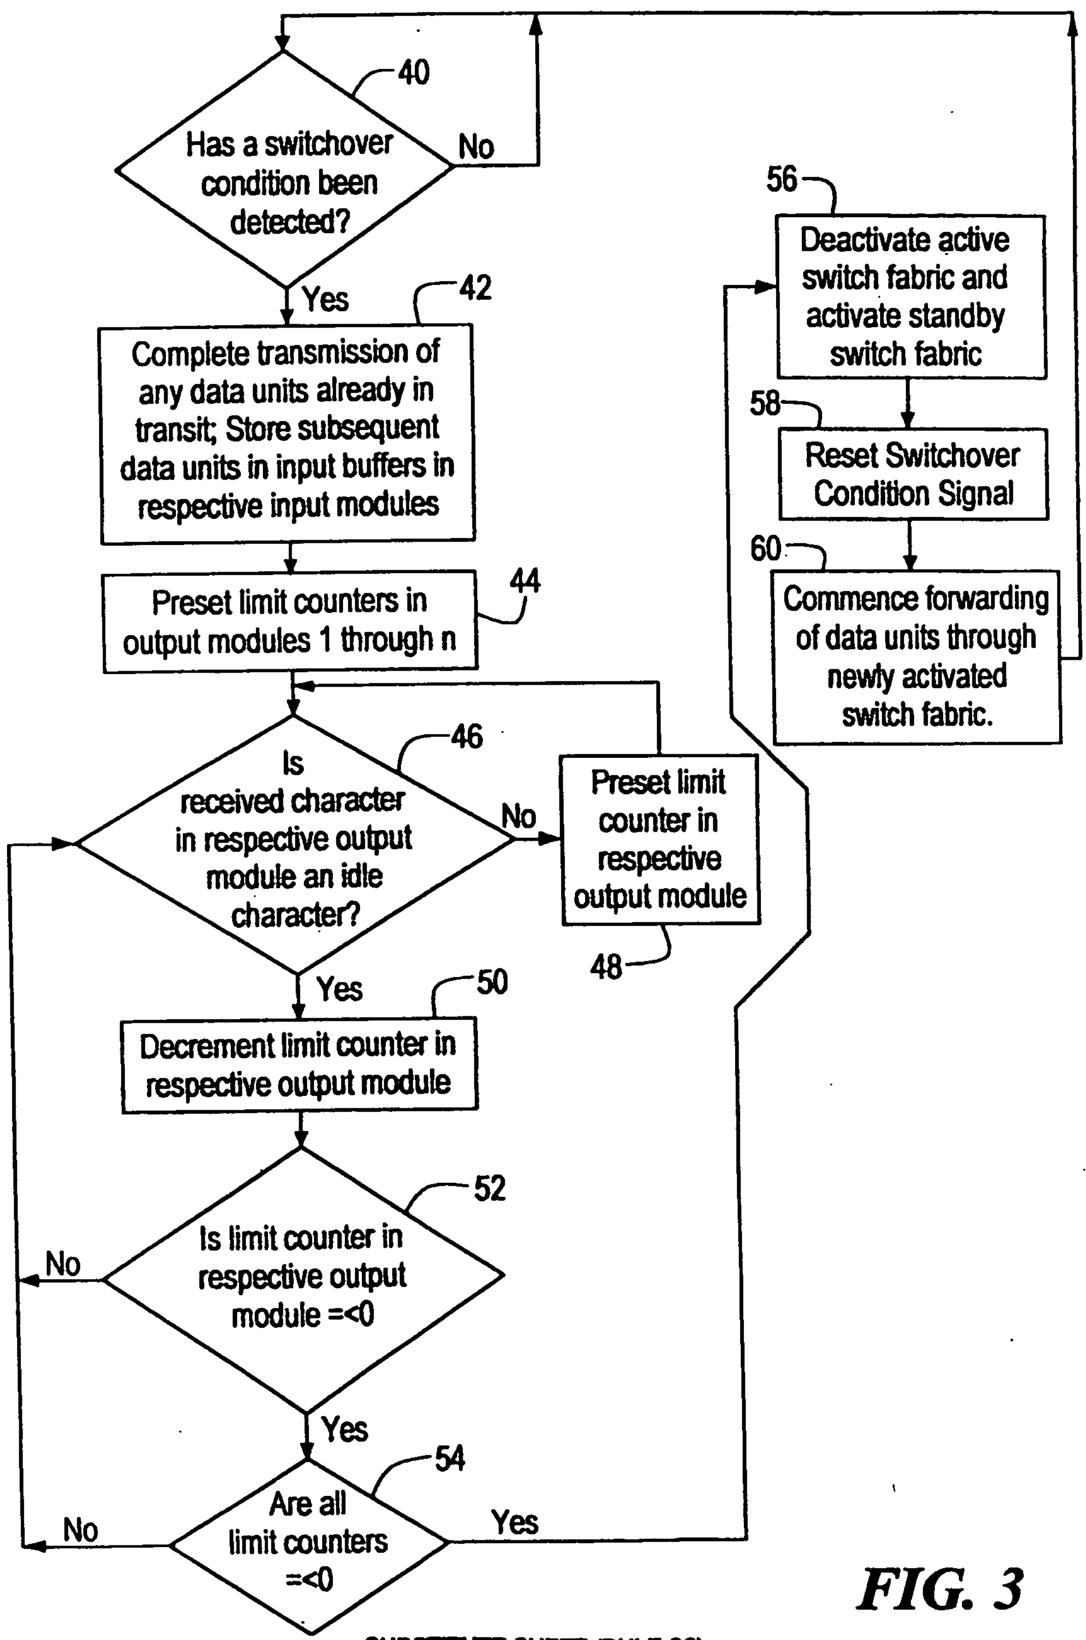

40
Has a switchover condition been detected?
No
Yes
42
Complete transmission of any data units already in transit; Store subsequent data units in input buffers in respective input modules
44
Preset limit counters in output modules 1 through n
46
Is received character in respective output module an idle character?
No
Preset limit counter in respective output module
48
Yes
50
Decrement limit counter in respective output module
52
Is limit counter in respective output module =<0
No
Yes
54
Are all limit counters =<0
No
Yes
56
Deactivate active switch fabric and activate standby switch fabric
58
Reset Switchover Condition Signal
60
Commence forwarding of data units through newly activated switch fabric.

FIG. 3